McKinsey & Company

Risk & Risilience Practice

The AI-enabled risk function of the future

AI is reshaping how banks manage risk. The next challenge is scaling AI capabilities through a redesigned risk operating model.

This article is a collaborative effort by Andreas Kremer, Anke Raufuss, Sebastian Schneider, and Thomas Poppensieker, with Maarten Vermeyen, Marc Mitrovic, and Stephan Beitz, representing views from McKinsey's Risk & Resilience Practice.



November 2026

Banks have spent the best part of two decades investing in risk function governance, capabilities, and data infrastructure. That process, combined with a complex regulatory environment, has led to spiraling demands on resources, which many institutions have addressed through a combination of prioritization, simplification, and targeted automation. Yet the pressures on the risk function continue to mount.

AI offers a new route to more efficient and effective risk management, enabling faster analysis, better decision-making, and improved economics across many workflows. Many institutions are testing the technology on individual processes and seeing promising results. The next step for chief risk officers (CROs) is to identify which capabilities can be scaled across the risk function and determine how to avoid creating new risks in the process.

Building on our earlier work on the future of risk[1] and our recent engagements, this article drills down into the most promising AI-led innovations in risk management, their implications for operating models, and how the risk function can fulfill its unchanged purpose in fundamentally new ways.

From manual oversight to continuous risk intelligence

The practice of risk management combines deep analysis of structured and unstructured data, continuous monitoring of regulation, use of quantitative models, and exercise of expert judgment. CROs now have an opportunity to redesign selected end-to-end processes around these capabilities.[2] With appropriate validation and oversight, AI-enabled workflows can monitor exposures more frequently, interpreting inputs, delivering more timely analysis, compiling evidence, and escalating exceptions to specialists. This higher level of automation enables risk managers to assign more resources to emerging threats and create value through actions such as challenging consequential decisions.

As institutions become more accustomed to AI's capabilities, some are using the technology across functional applications. Increasingly, we see integrated AI-enabled risk and finance platforms that combine information retrieval, quantitative analysis, quality checks, and reporting. For institutions with reliable data and effective controls, connecting balance sheet data, approved models, and market signals can shorten cycles of scenario analysis.

[1] Anke Raufuss, Arvind Govindarajan, Ida Kristensen, and Thomas Kelepouris, "The future of risk: How global trends are reshaping risk management," McKinsey, December 17, 2025.
[2] Anke Raufuss, Arvind Govindarajan, Ida Kristensen, and Thomas Kelepouris, "The future of risk: How global trends are reshaping risk management," McKinsey, December 17, 2025.

It is easy to see how AI applications can benefit individual disciplines: Credit officers can receive pre-analyzed underwriting materials sooner, while enterprise risk teams can devote more time to concentration and emerging threats. The question is how these changes affect the risk function as a whole. We see five structural shifts at play:

1. *Risk management becomes predictive, explorative, and focused on foresight.* Early-warning indicators, anomaly detection, and automated scenario modeling synthesize a broader set of internal and external signals at greater depth and speed. Deteriorating credit or operational conditions can be identified before they manifest as formal covenant breaches or credit losses, enabling earlier intervention and more proactive portfolio steering.

2. *Risk management becomes dynamic.* AI enables continuous, event-triggered monitoring of borrower positions, financial covenants, risk limits, policy adherence, model drift, and risk appetite. Automated quality checks, evidence logging, and exception-based routing increasingly support more responsive and targeted risk assessments.

3. *Risk management becomes faster and more consistent.* Agentic workflows accelerate end-to-end processes by automating data extraction, analysis, documentation, and routine policy checks. Specialized quality assurance capabilities help verify outputs against internal policies and approved regulatory requirements, while automated execution logs strengthen traceability. Human experts remain accountable for complex judgments, exceptions, and consequential decisions.

4. *Risk controls become embedded in end-to-end processes.* Automated controls and evidence capture are integrated directly into workflows spanning business units and infrastructure functions. This allows second-line risk teams to spend less time on manual oversight activities and concentrate on independent challenges, strategic policy setting, and systemic assurance, while strengthening transparency across governance lines. It also allows the first line to make faster decisions and provide better advice to their clients.

5. *The operating model shifts toward reusable analytical platforms.* Instead of managing static, siloed workflows, institutions deploy centralized platforms housing modular capabilities for credit intelligence, operational-risk sensing, model life cycle engineering, enterprise stress testing, and AI governance. This modularity industrializes execution while keeping ultimate accountability with human decision-makers.

Together, these dynamics represent a significant competitive shift in risk capabilities for banks. While banks are currently focused on using AI to increase productivity and enable them to adopt new regulation more efficiently, the real competitive opportunity comes through better foresight and business enablement. This can take the form of much more dynamic portfolio strategies and lending policies; better understanding of concentrations and transmission channels for risks; faster reactions to sector trends and disruptions; and deeper client discussions on balance sheet structuring, lending, and investment decisions.

Our recent client work has focused on AI's impact across sectors, especially as market volatility pushes companies along the AI value chain to prove that productivity gains and new business opportunities justify their investments. The scenarios detailed in this article affect not only private credit flowing into data centers but also the speed of AI adoption across industries, credit spreads and interest rate levels, private investments, and multiple parts of lending activities.

For CROs, the priorities are to select investments tied to risk outcomes, establish reusable capabilities, and clarify governance and talent requirements.

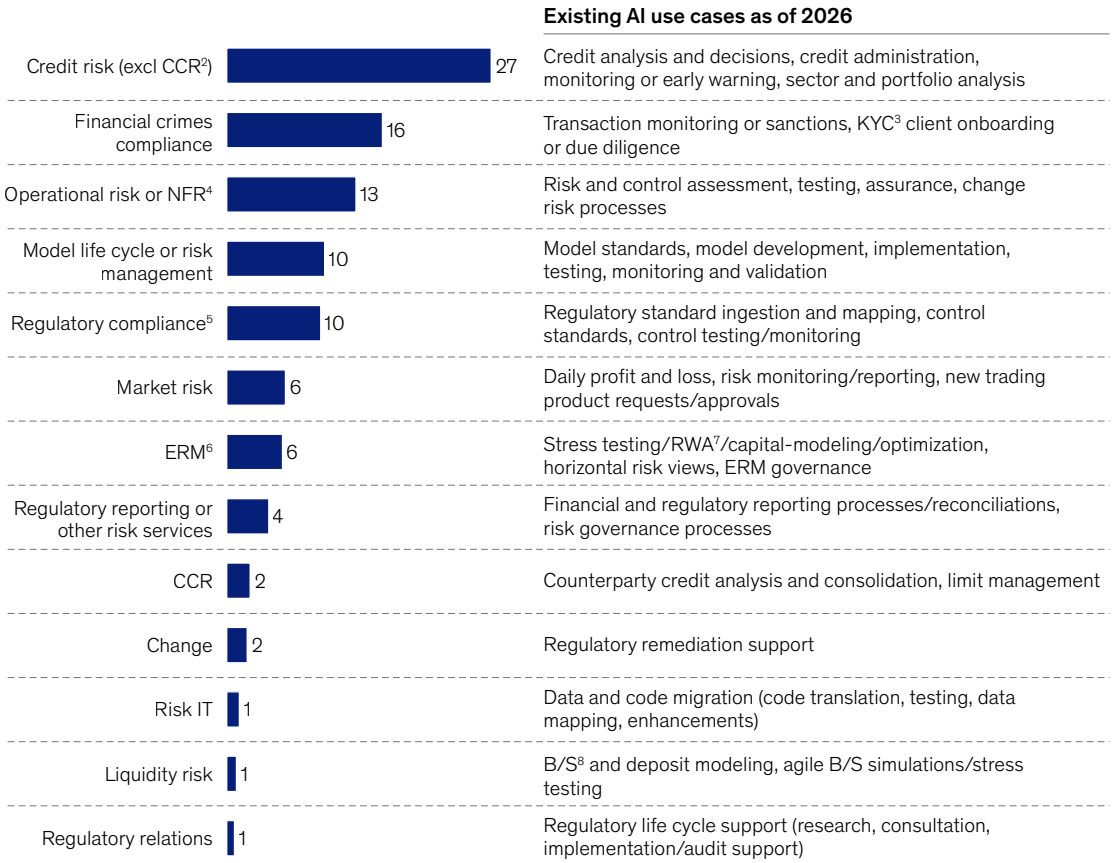
Where banks are deploying AI

Banks are commonly deploying AI across five risk domains (Exhibit 1). Below we highlight key learnings and remaining challenges.

Exhibit 1

Banks are already implementing or industrializing AI-enabled workflows across core risk processes.

Typical distribution of second-line-of-defense[1] risk or compliance full-time equivalents by service, large international or global banks, % of total



Note: Figures do not sum to 100%, because of rounding.
[1]Refers to the independent risk management and compliance functions that oversee and challenge the first line (business units, sales, and product teams), which owns and manages risk day to day. [2]Counterparty credit risk. [3]Know your customer. [4]Nonfinancial risk. Operational risk and NFR include areas such as operational resilience, third-party, technology/cyber, and conduct risk. [5]Regulatory compliance outside financial crime; eg, conduct, consumer protection, and regulatory change management. [6]Enterprise risk management. [7]Risk-weighted assets. [8]Balance sheet.
Source: McKinsey Global Risk Benchmark; QuantumBlack, AI by McKinsey; McKinsey analysis

McKinsey & Company

Wholesale lending: Integrating underwriting and continuous borrower monitoring

Banks are testing agentic workflows across wholesale lending, including retrieval of borrower data, financial analysis, credit policy checks, documentation, and exception management. Early implementations show productivity improvements of 30 to 50 percent across the end-to-end credit life cycle as well as shorter turnaround times and output quality rates of approximately 95 percent. Credit officers continue to review complex cases, policy exceptions, and final credit decisions.[3]

The degree of automation is calibrated to transaction complexity and exposure. Standardized, lower-risk transactions are handled through highly automated workflows, while moderately risky transactions retain human-assisted processes. Complex and high-exposure transactions are dominated by human-led assessments supported by AI-generated analysis. Leading institutions are extending automated credit assessment to exposures exceeding €5 million, with credit officers retaining accountability for final approvals.[4]

Banks are also beginning to connect credit origination with ongoing borrower monitoring. Following underwriting, AI-enabled workflows can track borrower filings, rating changes, market developments, and sector-specific risks, triggering targeted reassessments of debt service capacity, covenant headroom, and facility limits. Where integrated enterprise data and analytical capabilities are available, credit teams can also assess the implications of macroeconomic and sector-specific scenarios across individual borrowers and portfolios. The practical benefit is earlier identification of borrower deterioration and more timely intervention.

Model development and validation: Accelerating the model life cycle

AI is increasingly supporting quantitative model development across requirements interpretation, data preparation, code generation, benchmarking, testing, documentation, and legacy-code migration. AI-enabled tools can identify gaps in system requirements, generate test scripts, refactor legacy code, and support back-testing in controlled environments.

Early implementations show productivity improvements of 20 to 40 percent across model development activities, potentially reducing the time required for complex development cycles from months to weeks and accelerating routine recalibrations and documentation updates. These efficiencies allow quantitative teams to devote more attention to model conceptualization, statistical assumptions, and business applicability.[5]

AI is also advancing model validation through automated dossier completeness checks, independent code reviews, analytical replication, benchmarking, test generation, and regulatory evidence preparation. A reported reduction in validation time of 30 to 50 percent enables validators to expand testing coverage and focus more closely on model limitations, methodological soundness, and material risks. Effective implementation requires clear separation of development and validation responsibilities, approved testing standards, and independent human challenge.[6]

[3] "The future is agentic: AI's role in the end-to-end corporate credit process," McKinsey, December 12, 2025. Data based on McKinsey's most recent client engagements, with productivity improvements potentially reaching 80 percent across the end-to-end wholesale credit workflow.

[4] "Extracting value from AI in banking: Rewiring the enterprise," McKinsey, December 9, 2024. Additional data from McKinsey analysis and client engagements.

[5] Simona Viscardi, Carlos Díaz, and Olivier Maillet, "Evolving model risk management in the age of AI," McKinsey, July 22, 2026.

[6] Simona Viscardi, Carlos Díaz, and Olivier Maillet, "Evolving model risk management in the age of AI," McKinsey, July 22, 2026.

Nonfinancial and operational risk: Connecting risk signals and control assessments
AI is enabling continuous assessment of operational risks and controls by connecting information across business processes, policies, operational metrics, incident logs, audit findings, and external threat indicators. This is particularly relevant for risk and control self-assessments, which traditionally require substantial manual coordination, and periodic assessments across business units.

AI agents can prepopulate assessments, identify inconsistencies, evaluate controls against evolving regulatory requirements, and recommend adjustments following changes in business processes. By connecting these activities with operational data and incident information, institutions can identify emerging control weaknesses, detect patterns across business units, and provide more timely insights into process owners.

Another important development is the use of AI to maintain risk taxonomies, control libraries, and regulatory mappings. AI-enabled workflows can analyze internal policies and regulatory texts, identify affected controls, and prepare proposed updates for human review and approval. Early implementations show that selected taxonomy and control-mapping activities can be completed in three to four weeks, potentially shortening processes that previously required several months.[7]

Second-line teams use the resulting information to identify control weaknesses and prioritize mitigation. Over time, increasingly connected risk and control data provide the foundation for operational digital twins that simulate the effects of process changes, disruptions, and control failures.

Financial-crime compliance: Deeper customer analysis and more reliable transaction monitoring
Agentic AI can combine public records, corporate ownership registries, customer information, and transaction data to support anti–money laundering, know-your-customer (KYC), sanctions screening, and transaction monitoring. These workflows help investigators assemble relevant evidence and assess complex relationships, while authorized professionals retain responsibility for decisions.

In KYC, AI-enabled workflows track changes in ownership registries, sanctions and politically exposed persons lists, and adverse media exposures to flag cases for reassessment. They can also support document checks, beneficial-ownership analysis, screening, and requests for missing information. Early implementations in ultimate beneficial-ownership determination have led to handling-time reductions of more than half. Analysts then focus on complex ownership structures, conflicting information, and higher-risk relationships.

In transaction monitoring, AI-enabled workflows support alert triage, investigation documentation, and preparation of suspicious-activity reports (SARs). These applications offer opportunities to reduce investigation effort and improve consistency, subject to validation and regulatory oversight. False positives typically account for more than 95 percent of alerts in some transaction-monitoring environments, highlighting the potential value of more precise triage. Institutions are also exploring how investigator decisions and case outcomes can refine detection patterns, with changes subject to testing and approval.[8]

[7] Anke Raufuss, Björn Nilsson, Laura Webanck, and Andrew Margrave, "The bank's digital twin: The future of nonfinancial risk management," McKinsey, October 2, 2026.
[8] Based on McKinsey analysis and recent client engagements.

Enterprise risk management: Enabling integrated scenario analysis and strategic decision support
Enterprise risk management (ERM) offers a way to connect advances in analytics with decisions on portfolio positioning, capital, funding, and liquidity. By bringing together information across credit, market, liquidity, and interest-rate risk, these platforms enable teams to evaluate portfolio sensitivities under dynamic macroeconomic scenarios and idiosyncratic stress events.

Institutions with sufficiently integrated data and analytical capabilities can assess the combined capital, funding, and liquidity implications of unfolding market developments within 24 hours, supporting faster evaluation of potential management responses.[9]

An important enabler is AI-supported data reconciliation across ledger, risk, and trading systems. By aligning aggregated enterprise metrics with granular balance sheet positions, institutions can improve the consistency and timeliness of risk information. Integrated analytical platforms can then translate macroeconomic scenarios into portfolio-level impacts, identify potential risk thresholds, and evaluate the implications of balance sheet and hedging decisions.

For risk leaders, the question is how to organize data, analytical capabilities, and decision rights so that faster scenario analysis leads to timely management action.

What this means for the risk operating model

As banks roll out AI across domains, they must make changes to the risk function's capacity, decision-making processes, delivery model, skills, and architecture.

Accelerate elimination of low-value, labor-intensive tasks
The automation of manual data gathering, spreadsheet cross-checking, and narrative documentation creates opportunities for structural efficiency gains of 30 to 60 percent across selected risk workflows (net of human validation and oversight).[10]

Higher levels of efficiency enable institutions to reassess their nearshore and offshore captive delivery models. While labor arbitrage historically lowered unit costs, AI-enabled workflows offer a different cost and delivery model, with potential advantages in processing speed, consistency, and scalability, particularly when accounting for high operational attrition and onboarding costs. Some leading banks are targeting structural productivity gains significantly above 30 percent across their risk functions over the next two to three years, compared with the 5 to 10 percent annual targets common in traditional productivity approaches.[11]

Reorient toward strategic foresight and faster decision-making
Risk teams can use more timely analysis to challenge portfolio assumptions, revise risk appetite, and support management decisions.

With more timely and consistent data, risk teams can identify emerging developments earlier and adjust portfolio strategies and underwriting guidelines more quickly. For commercial banks, this capability translates into opportunities for competitive advantage: They can dynamically adjust sector risk appetites,

[9] Anke Raufuss, Arvind Govindarajan, Ida Kristensen, and Thomas Kelepouris, "The future of risk: How global trends are reshaping risk management," McKinsey, December 17, 2025. Data based on McKinsey analysis and recent client engagements.
[10] "Global risk productivity survey: Four themes shaping risk management," McKinsey, February 17, 2026. Data based on McKinsey analysis and recent client engagements.
[11] "How banks can boost productivity through simplification at scale," McKinsey, March 13, 2025. Data based on McKinsey analysis and recent client engagements.

reprice portfolios in response to changing market conditions, and proactively engage clients during periods of macroeconomic stress with custom financing solutions.

These capabilities also reinforce the role of a strategic ERM function operating alongside specialized portfolio teams. Risk professionals will need to strengthen their capacity for independent challenges, critical thinking, and strategic analysis to translate AI-generated insights into effective decisions.

Build a data-driven, agile function for continuous adaptation

Delivering AI-driven benefits requires risk functions to adapt processes and analytical capabilities faster. Multiagent networks and AI workflow factories, which can comprise 50 to 150 specialized agents, accelerate the development and modification of risk workflows. Minimal viable products (MVPs) for sophisticated credit or compliance workflows can be developed in three to four months, with functional expansions potentially delivered within weeks.

Traditional delivery models often delay critical process and policy changes as they compete for priority in IT backlogs. In an agentic architecture, policy updates and workflow instructions are supported by human-readable operational prompts and modular rule definitions. AI-assisted code generation, prompt engineering, and test validation accelerate implementation, supported by controlled testing, human approval, and appropriate technology governance.

Create an AI-fluent risk organization

Greater organizational agility depends on an AI-fluent risk workforce. Risk professionals across the organization need to understand agentic-system capabilities, limitations, and failure modes, alongside the domain expertise required to challenge their outputs and exercise sound judgment.

Risk teams must manage AI agents with disciplined performance oversight and clear accountability. This includes monitoring for model and workflow drift, implementing quality control loops, and driving continuous capability refinement. Because AI workflow instructions increasingly encode elements of risk procedures, risk functions need embedded technical and prompt-engineering experts, who should work closely with enterprise technology teams while retaining ownership of risk-specific requirements and decisions.

Scale the architectural and platform thinking

Scaling these capabilities requires a coordinated approach to technology architecture and platform governance. Uncoordinated point solutions can create fragmentation, technical debt, and rising operating costs. Leading institutions are building disciplined platform architectures backed by a shared data and analytical infrastructure, common interfaces, modular reasoning capabilities, and centralized governance standards. The architecture supports faster workflow development, interoperability, and more efficient reuse of AI capabilities across risk domains.

Effective implementation also requires clear responsibilities for human oversight, agent interactions, and platform ownership, supported by governance frameworks that establish standards for output quality, factual accuracy, traceability, and regulatory compliance.

Deciding where responsibilities sit

Banks that make the structural adjustments discussed above will create a model combining enterprise-wide risk oversight, specialist risk judgment, and shared AI capabilities. Within that architecture, banks will need to decide where responsibilities and decision rights sit. One possible approach consists of three complementary pillars:

— *A strategic, horizontal risk function* steers enterprise-wide resilience, capital allocation, macroeconomic risk assessment, and portfolio positioning, integrating insights across risk disciplines to support senior management decisions.

— *Lean, AI-enabled risk domains* span credit, market, operational risk, and compliance, with professionals focused on complex underwriting, strategic policy challenge, exception management, and independent oversight of AI-supported decisions.

— *A center of competence for risk AI* develops and maintains reusable AI capabilities, supports analytical platforms, establishes engineering standards, and monitors model and workflow reliability in collaboration with enterprise technology and governance functions.

In this three-tiered model, the horizontal risk function sets enterprise priorities and provides independent challenge. Specialized risk domains retain accountability for risk decisions, with shared capabilities under common engineering and governance standards. Clear decision rights across these groups help avoid fragmented ownership and preserve effective oversight (Exhibit 2).

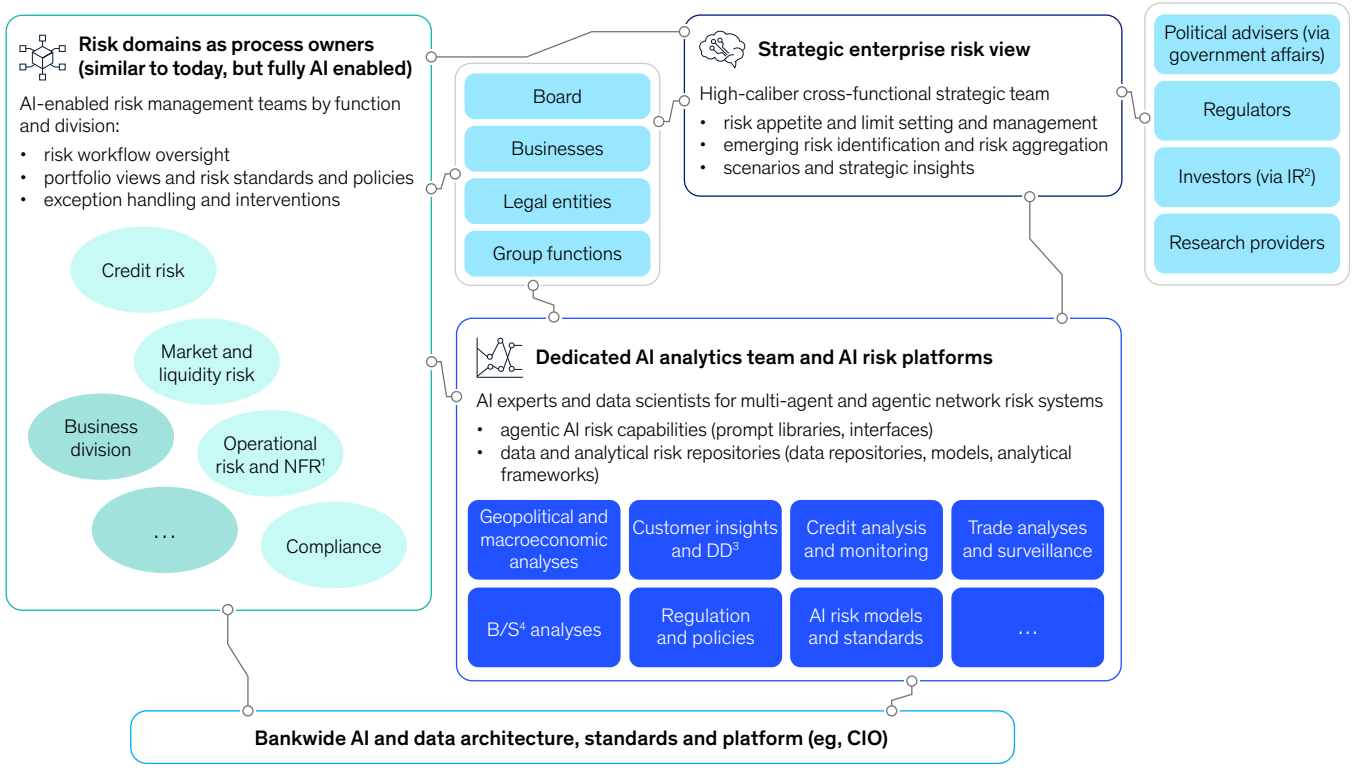
The transition also creates opportunities to fundamentally reallocate capacity within the risk function. We estimate that AI-enabled workflows could eliminate 60 to 70 percent of manual operational tasks, depending on the process and degree of automation. Institutions could reinvest half of the released capacity into deeper analytical capabilities, strategic risk advisory, and faster portfolio steering, while directing the remaining capacity toward specialized AI engineering, data governance, and platform development.[12]

The strategic value depends on redeploying released capacity to strengthen risk analysis, controls, and decision support.

[12] *Global Banking Annual Review 2025: Why precision, not heft, defines the future of banking*, McKinsey, October 23, 2025. Data based on McKinsey analysis and recent client engagements.

Exhibit 2

An AI-enabled risk operating model provides AI competence for risk functions.



[1]Nonfinancial risk.
[2]Investor relations.
[3]Due diligence.
[4]Balance sheet.

McKinsey & Company

Addressing five potential roadblocks to scaling

Implementing a new model requires changes to workforce planning, technology delivery, governance, and investment priorities. Five potential roadblocks deserve attention:

1. *Underestimating organizational change.* Traditional cost programs often target 10 to 20 percent head count rationalization, while nearshore and offshore delivery models relocate activities to lower-cost locations. AI-enabled transformation can eliminate entire categories of transactional tasks, fundamentally changing workflows, staffing requirements, and oversight responsibilities. Risk leaders must redesign processes, role profiles, and workforce capacity alongside AI deployment, supported by deliberate change management and plans to redeploy talent into higher-value activities.[13]

2. *Deficits in embedded domain AI talent.* Many implementations rely on small, centralized teams of external consultants or corporate data scientists whose involvement diminishes after initial deployment. Sustaining and improving AI-enabled workflows requires embedded technical expertise within risk domains, combined with a deep understanding of risk processes and decision requirements. Risk functions need talent strategies that integrate technical upskilling for risk professionals, targeted recruitment of AI and data engineering specialists, and close collaboration with enterprise technology teams.

[13] "Global risk productivity survey: Four themes shaping risk management," McKinsey, February 17, 2026. Data based on McKinsey analysis and recent client engagements.

3. *IT governance friction and inefficient economics.* Legacy technology delivery models can slow AI deployment, while inefficient model selection and fragmented architectures increase operating costs. Institutions can improve implementation speed and economics by establishing modular agentic workflows, using appropriately sized and specialized models, and streamlining development and approval processes within established governance standards. Illustrative implementation estimates suggest computational token costs of $30,000 to $70,000 for developing an automated nonfinancial risk framework that processes hundreds of policies and maps risk taxonomies. Estimated inference costs for individual workflows range from $100 to $250 for a complex corporate credit memo review and $20 to $30 for a transaction-monitoring case review. These figures highlight the potential economics of AI-enabled execution, although total costs must also account for integration, validation, infrastructure, and ongoing monitoring.[14]

4. *Absence of clear AI risk governance standards.* Unclear approval requirements for generative and agentic systems can delay production deployment and leave promising applications confined to pilot environments. Effective scaling requires senior-level endorsement of governance standards, jointly supported by risk and technology leadership. These should establish clear accountability, risk-based approval processes, and ongoing performance oversight. Four operational criteria are important:

    - **Completeness:** Confirming that mandatory factual, analytical, and regulatory elements have been addressed

    - **Accuracy:** Verifying that outputs and conclusions are supported by approved data, models, and source information

    - **Replicability:** Assessing the consistency and reliability of outputs across repeated executions and defined testing conditions

    - **Traceability:** Maintaining appropriate records of inputs, sources, system actions, decisions, and human interventions to support auditability and independent review

5. *Fragmented, bottom-up transformation strategies.* Disconnected proofs of concept can demonstrate technical potential without materially changing the risk function's operating baseline. Overcoming this "lighthouse trap" requires senior risk leadership to establish an enterprise-wide transformation road map, prioritize use cases with significant productivity and decision-making benefits, and invest in reusable data, technology, and governance capabilities. Progress should be measured against improvements in end-to-end processes, risk outcomes, and organizational capacity, ensuring that individual deployments contribute to the broader operating-model transformation.

Overcoming these barriers requires CRO sponsorship and measurable objectives. Leading banks track end-to-end process performance, decision quality, control effectiveness, and the redeployment of capacity.

[14] "Extracting value from AI in banking: Rewiring the enterprise," McKinsey, December 9, 2024.

Find more content like this on the
McKinsey Insights App
Scan • Download • Personalize

Moving to AI-enabled risk insights as a key competitive advantage

Artificial intelligence represents a centennial shift for knowledge-intensive, analytical industries, especially for banking risk functions historically constrained by manual reviews, administrative overhead, and dense regulatory reporting.

The transition to an AI-enabled risk operating model is not a speculative question of *if*, but an operational mandate of *when* and by *whom*. Early adopters will capture substantial cost and productivity advantages and will establish an enduring competitive edge through faster, more sophisticated decision-making. In an increasingly volatile operating environment defined by sudden macroeconomic, geopolitical, and technological shocks, decision speed and portfolio foresight are the ultimate drivers of institutional resilience. The cost of delay will vastly exceed the investment required to lead.

Andreas Kremer is a partner in McKinsey's Berlin office, **Anke Raufuss** is a senior partner in the Sydney office, **Sebastian Schneider** and **Thomas Poppensieker** are senior partners in the Munich office, **Maarten Vermeyen** is an expert in the Brussels office, and **Marc Mitrovic** is a senior expert in the Frankfurt office, where **Stephan Beitz** is an associate partner.

Copyright © 2026 McKinsey & Company. All rights reserved.